

(12) UK Patent Application (19) GB (11) 2 346 280 (13) A

(43) Date of A Publication 02.08.2000

(21) Application No 9823015.4

(22) Date of Filing 22.10.1998

(71) Applicant(s)
Hewlett-Packard Company
(Incorporated in USA - Delaware)
3000 Hanover Street, Palo Alto, California 94304,
United States of America

(72) Inventor(s)
Michael Peter Spratt

(74) Agent and/or Address for Service
Agilent Technologies UK Ltd
Legal Dept, Eskdale Road, Winnersh Triangle,
WOKINGHAM, Berks, RG41 5DZ, United Kingdom

(51) INT CL⁷
H04Q 11/00, H04J 14/02

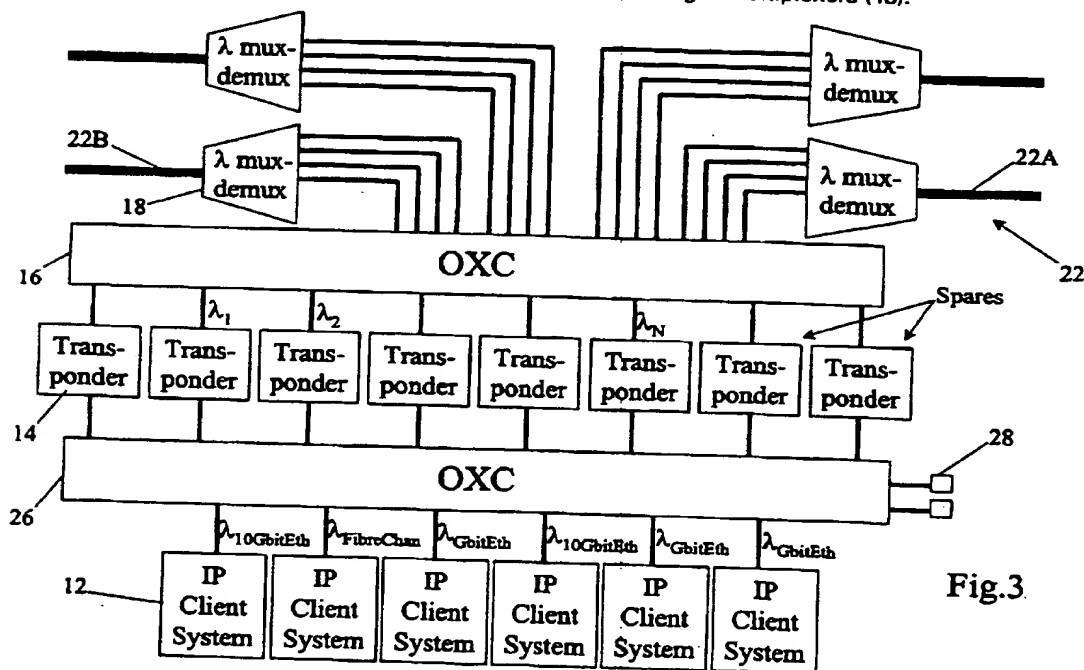
(52) UK CL (Edition R)
H4B BN

(56) Documents Cited
WO 97/21289 A1 US 5303077 A


(58) Field of Search
UK CL (Edition Q) H4B B1X BN
INT CL⁶ H04J 14/02, H04Q 11/00
ONLINE: WPI, EPODOC, JAPIO

(54) Abstract Title
Optical switching interface using transponders

(57) In a optical switching interface trunk fibres (22) are switchably coupled to client systems (12) via two optical cross-connects (OXC - 16, 24) interconnected by a set of transponders (14) performing wavelength conversion and other functions as desired. The inclusion of the OXC (24) coupling the client systems to the transponders provides useful additional functions and flexibility in the switching interface. The first OXC (16) may be connected to the trunk fibres (22) via one or more wavelength multiplexers (18).



GB 2 346 280 A






Fig.2




Fig.3

Optical switching interfaces

[30980151 EP]

Technical Field

5 This invention relates to optical switching interfaces for selectively coupling optical signals between end-points of an optical communication system and an optical link.

Background Art

10 It is anticipated that new optical telecommunications networks will be implemented in metropolitan areas, and one possible interface between the telecommunications network and the end customer will be "Gigabit Ethernet" (or a higher speed derivative). In this scenario, the end customer's equipment (for example a router) will have a Gigabit Ethernet interface card, which may for example connect to the telecommunications network via a "transponder" - an optical frequency converter for converting an optical signal between the end-user system wavelength and the wavelength of an assigned optical "channel" in a wavelength division multiplexing (WDM) system. (In an all-optical network an optical channel can be considered to exist, constituting a communications path between two network endpoints; although signals traversing this channel may in reality be optically amplified and may pass through optical cross-connects, wavelength multiplexers and demultiplexers, and wavelength converters, the channel is still in principle viewed as a single integral end-to-end path.) In initial implementations it is likely that an optical channel will be set up across the optical telecommunications network to a second corresponding transponder and then on to a second end customer's equipment. It will therefore create a point to point connection, across the metropolitan area telecommunications network, between the two pieces of customer's equipment.

25 WDM technology and the use of optical cross-connects (OXC - see for example US patents 5,699,462 and 5,732,168) are now being implemented in working products. As well as the use of an optical layer as the highest layer in Synchronous Optical Network (SONET)/Synchronous Digital Hierarchy (SDH) technology, there is also very considerable interest in Internet Protocol (IP) and Asynchronous Transfer Mode (ATM) networks being implemented in a way which minimises or eliminates SONET/SDH support within the protocol stack. Specifically, it seems that in the long-haul networks, because of the considerable number of electrical repeaters already in place which are specific to SONET or SDH, it will be required in the short term to keep SONET/SDH framing, but the SONET/SDH multiplexing and cross-connect functions can potentially be eliminated.

30 In the metropolitan / municipal area, there is no requirement for electrical repeaters, and so both IP/ATM within SONET/SDH framing, and independent IP or ATM over WDM are likely to arise. Of particular interest will be the use of Gigabit Ethernet technology to implement IP over the WDM optical layer channel.

A parallel development is the realisation that the tremendous current and projected future demand for bandwidth will comprise packetised data rather than voice communications. The demand for data is projected to considerably exceed that of voice within very few years. It therefore makes sense to optimise the telecommunications network for data traffic rather than voice. There is also a strong movement towards transporting voice across IP networks (IP telephony), and this leads to the conclusion that IP may become the universal convergence layer.

It is an object of this invention to provide methods and apparatus for optical switching and interfacing which facilitate implementation of optical communications systems in these circumstances.

Disclosure of Invention

According to one aspect of this invention there is provided a switching interface for selectively coupling optical signals between end-points of an optical communication system and an optical link, comprising: a plurality of transponders; a first optical cross-connect switch for selectively coupling said transponders to said optical link; and a second optical cross-connect switch for selectively coupling said transponders to said end-points.

Brief Description of Drawings

An optical switching interface in accordance with this invention will now be described, by way of example, with reference to the accompanying drawings, in which:

- Figure 1 shows a first switching interface;
- Figure 2 shows a second switching interface; and
- Figure 3 shows a third switching interface.

Best Mode for Carrying Out the Invention, & Industrial Applicability

Various kinds of client enterprise or organisation, and in their respective requirements for a metropolitan communications network, can be envisaged:

- (a) A client has a requirement for a permanent point to point optical channel to connect two sites. The client expects the optical channel to be highly reliable, i.e. protection/restoration by the metropolitan communications network is a requirement.
- (b) A client has a requirement for a permanent point to point optical channel to connect two sites. The client also has a lower bandwidth alternative route available between the two sites, and would consider the service satisfactory if any break in service on the optical channel were restored within a few hours.
- (c) A client requires point to point optical channel connections, but the client specifically wishes to handle protection themselves e.g. by having two diversely routed optical channels, and two (rather than one) Gigabit Ethernet cards at each end of the link.

(d) A client requires access to a point-to-point optical channel every evening for an hour, starting at a particular time, in order to do a major backup of the data (data bases, transaction logs etc.) which are underpinning the business. This major backup onto a completely separate site protects against a major loss of data (which could potentially be caused by fire, flood or other disasters).

(e) A client requires bandwidth on demand between a site and one of the other sites across the metropolitan communications network. The time and duration of the demand for the optical channel cannot be predicted in advance, and neither can the end points (i.e. if a client organisation has three sites, the two end points could be any two of the three sites). Potential applications are: large general demand for data transmission which outstrips the current connectivity between the sites; video conference; high quality digital video transmission; transferring large images (e.g. publishing industry, also healthcare); transfer of large amounts of engineering design data; and transferring large amounts of software.

This is not an exhaustive list, and it is also clear that administratively a variety of connections are possible, e.g. both endpoint sites belonging to the same organisation, connection from a site to an ISP, or connection between two separate organisations.

It follows that an IP-based metropolitan communications network may differ in various respects from a solely SONET/SDH-based metropolitan network.

1. Different levels of protection/restoration have to be offered to different clients

Some clients require optical protection, but some do not. In particular, the client in scenario (c) does not require protection/restoration. Also the client in scenario (e) probably does not want protection/restoration, in that the service being requested from the metropolitan communications network is best-effort, i.e. there is a possibility that a demand for an optical channel may have to be refused owing to lack of capacity. Therefore, very reliable service in terms of protection/restoration is probably not required by this client.

Protection does not come for free and is potentially quite a large added cost to the customer. There are costs associated with the extra optical channel required, the extra equipment, etc. Therefore clients not requiring protection/restoration should not be expected nonetheless to pay for it. From the operator's viewpoint, channel capacity freed up by not having to offer protection/restoration can be used to generate other business. It is therefore preferable to design the system so that protection/restoration is not mandatory on every optical channel.

As a corollary, this argues more in favour of protection/restoration at the wavelength level rather than the fibre level.

2. There has to be a mechanism for the client router to exchange management/signalling messages with the Element Manager in the Optical Cross-Connect System

In scenario (e), the client requiring bandwidth on demand will have to send a message to the Element Manager in the Optical Cross-Connect System, requesting an optical channel to a particular end point. Similarly a reply has to be sent from the Element Manager back to the client, indicating a successful conclusion, or that not enough capacity on the relevant route is currently available (an 'engaged tone message').

One cannot assume that the client router can send a message to some more global Network Management Station or Cell Controller, since one cannot guarantee that there will be any connectivity from the client router to this more global Network Management Station which does not go via the Optical Cross Connect System. (Also, from the IP protocol point of view, there is an issue as to how the client router would find the address of the correct global management entity.)

The management/signalling messages exchanged between the client router and the Element Manager will typically be in Simple Network Management Protocol (SNMP) format, and clearly the Optical Cross-Connect System needs to have (physical layer and Media Access Control layer) hardware which will send and receive management/signalling frames.

Management/signalling data could be sent at the same wavelength as the client IP frames, or could be sent on a separate optical supervisory channel. In the case of Gigabit Ethernet (and future higher speed derivatives) carrying the IP traffic, it seems that it would be possible and most cost effective for the management/signalling messages to be sent using the Gigabit Ethernet (or higher speed derivative) itself, and that the use of a separate supervisory channel would be unnecessary.

25

3. The demand for optical channels will be asymmetric in the IP Metro-Network case, as opposed to symmetric in the SONET case

This is particularly true in scenario (d) for example, where a major backup of business data will go from one site to a 'data warehouse' and little data (mainly higher-layer-protocol short-acknowledgement messages in fact) would go back from the 'data warehouse' to the client's site. It would therefore be more cost effective for the client to implement the low speed direction using some other telecommunications network, and just use an optical channel for the high speed direction. Therefore, in the IP metropolitan communications network case, the architecture involves unidirectional optical channels, whereas in the SONET case it could involve bi-directional channels.

35

4. In the IP metropolitan communications network case, there is a greater demand for wavelength converters (or other techniques), to counter the inefficiencies caused by

wavelength blocking

The environment in which the Optical Cross-Connect System is used is much more dynamic than the purely SONET environment. Optical channel connections will be set up and broken down by clients with considerable frequency, and also in a fairly random way.

- 5 It is therefore much more likely that the network will encounter wavelength blocking (i.e. a demand for a channel path occurs, and there is channel capacity on the path, but only at a wavelength the end systems cannot make use of). The end clients in the SONET case may also have very dynamic behaviour, but in the SONET case the optical layer is shielded by the SONET layer from the changes in the end clients' bandwidth demands.

10

5. In the IP metropolitan communications network case, it is possible to use the Cyclic Redundancy Check (CRC) of the Gigabit Ethernet (or higher speed derivative) frame to verify bit error rate performance

- 15 An assumption of the IP metropolitan communications network is that the predominant optical channels are carrying IP traffic. Inspecting whether a received CRC is the correct one provides a way to estimate the end-to-end bit error rate (BER) of the optical channel. There are a number of other checksums and CRCs which could be useful for checking the optical channel BER. Some examples are the Transmission Control Protocol (TCP) checksum, the Point-to-Point Protocol (PPP) CRC and the IP header
20 checksum. The BER measurement module could be made to switch between looking at the checksum/CRC of the various common protocols likely to be sent by the IP metropolitan communications network optical channel clients, the switch being initiated by network management messages. The BER measurement module could also be programmable at a lower level (also by special network management messages) to cope with future as well
25 as current protocols.

Various possible switching interface architectures which cater for one or more of these requirements are shown in Figures 1 to 3.

- Referring to Figure 1, the client systems 12 are coupled via transponders 14 (described in more detail below) to an OXC 16, which in turn is coupled via wavelength
30 multiplexer/demultiplexers 18 to a fibre switch 20 connected to the metropolitan network trunk fibre 22. Protection and restoration are carried out at the fibre level, by the fibre switch 20. This fibre switch is a smaller version of the OXC 16, and could be implemented with the same technology. Protection and restoration in this case means that if a fibre breaks, the signals sent in that fibre are re-routed to travel via a functional route.
35 For example, if in Figure 1 a fibre 22A shown extending to the right of the fibre switch 20 breaks, the signal is typically re-routed on a spare fibre 22B extending from the fibre switch 20 to the left.)

Wavelength converters 24 are included, connected to the OXC 16, so that if clients

require service, but wavelength blocking occurs, the wavelength converters 24 can convert to a wavelength which is available for use. However, this architecture does not allow different levels of protection/restoration for different clients.

Figure 2 shows an architecture in which all protection/restoration and switching are performed at the wavelength level. This architecture is similar to that shown in Figure 1, but the outputs of the multiplexer/demultiplexers 18 are connected directly to the trunk fibres 22 (i.e. the fibre switch 20 is omitted). All protection/restoration is handled on a per-wavelength basis, which allows the flexibility of enabling some clients to have their connection protected, while others can opt not to incur the cost overhead of such protection, or to handle the protection/restoration themselves. Where protection/restoration is implemented, it is appropriate for the BER or other relevant measurements to be carried out at the transponders 14, and for this information to be fed to the controller in the OXC 16 which can then make a fast decision to carry out protection/restoration switching.

Figure 3 shows a third architecture in which the outputs of the multiplexer/demultiplexers 18 are again connected directly to the trunk fibres 22, but in which a second OXC 26 is included, to provide flexible signal path switching between the client systems 12 and the transponders 14. This architecture has several advantages:

1. If clients use the optical network only intermittently (e.g. a backup for an hour each night) so that there is time sharing of the network by the clients, then there will be a considerable saving in transponder costs. Instead of one transponder 14 per client there is only a requirement of one transponder 14 per port of the OXC 26, which could be a much lower number.
2. The desirability of having on-line spare transponders is addressed. Upon transponder failure or fault, one of the spare transponders can be switched into action immediately. By contrast, in other architectures (e.g. those in Figures 1 and 2), there is no possibility of on-line real-time switching to a spare if a transponder 14 fails.
3. With the Figure 3 architecture, the requirement to have wavelength converters to prevent wavelength blocking is much reduced. Furthermore in the specific case where all clients use one protocol, such as Gigabit Ethernet, the requirement for wavelength converters is eliminated. In this case (ignoring spares) for each working fibre there is a Gigabit Ethernet transponder per wavelength i.e. there is a one-to-one correspondence between transponders and wavelengths. So if there is a spare wavelength available when a client requests a new optical channel, there must also be a transponder 14 available, so no wavelength blocking can occur.
4. There is also much greater flexibility, in that if a client changes for example from Gigabit Ethernet to a higher speed protocol, it is a simple matter to connect that client's system to the appropriate transponder 14.

5. Local clients can be directly interconnected without having to take up transponder resources.

6. It is also potentially possible for transit (bypass) traffic (e.g. a channel coming in on one the fibre 22A, and going out through the fibre 22B without going to a local client) to be actively repeated by going through back to back transponders 14, thus boosting the optical power level of the signal and removing/reducing distortions. This is accomplished by routing a signal coming in on the fibre 22A to pass through a wavelength demultiplexer 18 and the OXC 16, one of the transponders 14, and the second OXC 26, and then loop back through another transponder 14, the OXC 16 again and then out through a wavelength multiplexer 18 on to the fibre 22B.

7. The second OXC 26 can route management and signalling messages from the client systems 12 to associated management hardware 28.

As noted above, the transponders 14 are used to translate client signals (e.g. from routers, SONET modules, ATM switches etc.) from loosely defined and unstable wavelengths to well defined and stable wavelengths as they are transferred into a WDM system. Transponders also have some other advantages:

- they can act as a demarcation device between client and carrier, with 'keep alive' signal and test capabilities; as explained above, incorporation of BER measurement capabilities etc. is also possible;
- they transmit continuously, and so optical amplifier disruption, caused by varying total power levels due to client equipment turning on and off, is avoided;
- they ensure a consistent power level enters the optical network and so increase the margins/reach of the WDM network; also, a remote client becomes allowable;
- transponders could also be convenient places from an architectural viewpoint to locate various other functions, besides BER measurement and verification, such as:
 - channel encryption;
 - Forward Error Correction (FEC);
 - optical power measurement;
 - line code conversion (e.g. between different block codes);
 - frame format conversion (e.g. from Gigabit Ethernet to some other format);
 - channel protection (i.e. duplication of optical paths so that a spare path is available in case the working path fails);
 - network management (e.g. using SNMP or telecom management network (TMN)).

CLAIMS

[30980151 EP]

1. A switching interface for selectively coupling optical signals between end-points of an optical communication system and an optical link, comprising:
 - 5 a plurality of transponders;
 - a first optical cross-connect switch for selectively coupling said transponders to said optical link; and
 - a second optical cross-connect switch for selectively coupling said transponders to said end-points.
- 10 2. The interface of claim 1, wherein said first optical cross-connect is connected to said optical link via at least one wavelength multiplexer/de-multiplexer.



Application No: GB 9823015.4
Claims searched: Both

Examiner: Stephen Brown
Date of search: 19 May 1999

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK CI (Ed.Q): H4B (BN, BKX)

Int CI (Ed.6): H04J: 14/02, H04Q: 11/00.

Other: Online: WPI, EPODOC, JAPIO.

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X	WO 97/21289 A1 (Bell) See especially the abstract & figure 9.	1
X	US 5 303 077 (Standard Elektrik) See especially the abstract & figure 1.	1

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

THIS PAGE BLANK (USPTO)

THIS PAGE BLANK (USPTO)